BIOINFORMATICS

Vol. 00 no. 00 2015
Pages 1–9

# SANA: Simulated Annealing Network Alignment Applied to Biological Networks

Nil Mamano and Wayne Hayes*

Department of Computer Science, University of California, Irvine CA 92697-3435, USA

Received on XXXXX; revised on XXXXX; accepted on XXXXX

Associate Editor: XXXXXXX

arXiv:1607.02642v1 [q-bio.MN] 9 Jul 2016

**ABSTRACT**

The alignment of biological networks has the potential to teach us as much about biology and disease as has sequence alignment. Sequence alignment can be optimally solved in polynomial time. In contrast, network alignment is $NP$-hard, meaning optimal solutions are impossible to find, and the quality of found alignments depend strongly upon the algorithm used to create them. Every network alignment algorithm consists of two orthogonal components: first, an objective function or measure $M$ that is used to evaluate the quality of any proposed alignment; and second, a search algorithm used to explore the exponentially large set of possible alignments in an effort to find "good" ones according to $M$. Objective functions fall into many categories, including biological measures such as sequence similarity, as well as topological measures like graphlet similarity and edge coverage (possibly weighted). Algorithms to search the space of all possible alignments can be deterministic or stochastic, and many possibilities have been offered over the past decade. In this paper we introduce a new stochastic search algorithm called SANA: Simulated Annealing Network Aligner. We test it on several popular objective functions and demonstrate that it almost universally optimizes each one significantly better than existing search algorithms. Finally, we compare several topological objective functions using SANA. Software available at http://sana.ics.uci.edu.

**Contact:** whayes@uci.edu

## 1 INTRODUCTION

### 1.1 Context

Network alignment is the task of finding the best way to "fit" one network inside another. It has applications in several areas, including ontology matching (Li *et al.*, 2009), pattern recognition (Zaslavskiy *et al.*, 2009), language processing (Bayati *et al.*, 2009), and social networks (Zhang and Tang, 2013). Thus, the specific goal of network alignment depends on the context. We focus on a particular application from the computational biology domain: the alignment of protein-protein interaction (PPI) networks. In a PPI network, nodes represent proteins from a given organism, and edges connect proteins that interact physically. These kinds of interactions are discovered through high throughput experimental methods such as yeast two-hybrid screening (Ito *et al.*, 2000) or protein complex purification via mass-spectrometry (Krogan *et al.*, 2006).

*to whom correspondence should be addressed

PPI network alignment has many interesting applications. It can serve to transfer biological information across species (Kuchaiev *et al.*, 2010), which, in turn, has been used to offer insights on the mechanisms of human diseases (Uetz *et al.*, 2006), or the process of aging in humans (Milenković *et al.*, 2013).

Network alignment can be classified as local or global. The former aims to align small regions accurately (Kelley *et al.*, 2004). Consequently, it often fails to find large conserved connected subgraphs in different networks. By contrast, global network alignment aims to generate one-to-one node mappings between two networks. By aligning entire networks, it overcomes the shortcomings of local network alignment. For this reason, the majority of recent research has focused on global network alignment (see section 1.2). There are also methods that allow alignments between more than two networks (Liao *et al.*, 2009; Hu *et al.*, 2014).

We focus on pairwise global network alignment. Its goal is to find a one-to-one mapping from the proteins of the smaller network to proteins in the larger network. Ideally, we would like to find the most biologically relevant mapping: aligned proteins in both networks should be homologically related, in the sense that they used to be *same* protein in the species which was the common ancestor of the species of both PPI networks. Since proteins may have multiple descendants, such a mapping may not be one-to-one. For now we ignore this complication and view global pairwise 1-to-1 network alignment as a convenient approximation to the truth.

### 1.2 Previous work

Current PPI network alignment methods use a combination of biological and topological information to align similar proteins. Biological information includes *a priori* knowledge about the proteins, such as amino acid sequences. On the other hand, topological information is extracted exclusively from the structure of the PPI network. Since the goal is to obtain biologically relevant alignments, early methods focused on biological information. However, as our understanding of topology-function relationships (Davis *et al.*, 2015) has improved, topology information has gradually shifted to a central role. For instance, it has been shown recently that topological information is more important than sequence information for uncovering functionally conserved interactions (Malod-Dognin and Pržulj, 2015). Topological knowledge can be extracted in many forms. For example, the wiring patterns in the vicinity of homologous proteins in different networks tend to be similar. This information is well

© Oxford University Press 2015.

captured by graphlets (Pržulj *et al.*, 2004), which generalize the concept of the degree of a node.

There are a wide variety of network alignment methods. This diversity is motivated by the inherent computational complexity of network alignment: topologically speaking, one would like to find an alignment that maximizes the number of preserved interactions, i.e., interactions between proteins that are mapped to proteins that also interact. However, this problem is *NP-hard* because it is a generalization of *subgraph isomorphism*, which is *NP-complete* (Cook, 1971). This means that no efficient algorithm is known. Thus, practical methods must rely on approximation and heuristic techniques; and when it comes to heuristic algorithms, the possibilities are endless but there is no obvious "best option".

In the biological network domain, work over the past few years has included IsoRank (Singh *et al.*, 2008), the family of GRAAL algorithms (GRAAL (Kuchaiev *et al.*, 2010), H-GRAAL (Milenković *et al.*, 2010), C-GRAAL (Memisevic and Pržulj, 2012), MI-GRAAL (Kuchaiev and Pržulj, 2011), L-GRAAL (Malod-Dognin and Pržulj, 2015)), NATALIE (Klau, 2009), GHOST (Patro and Kingsford, 2012), NETAL (Neyshabur *et al.*, 2013), SPINAL (Aladağ and Erten, 2013), PISwap (Chindelevitch *et al.*, 2013) MAGNA (Saraph and Milenkovi, 2014) and its successor MAGNA++ (Vijayan *et al.*, 2015), HubAlign (Hashemifar and Xu, 2014), and OptNetalign (Clark and Kalita, 2015). Elmsallati *et al.* (2015) did a survey of existing methods, datasets, and optimization measures.

In general, each method defines an *objective function* or *measure* over alignments that can be viewed as a *score*, and then proposes a *search algorithm* that searches through the enormous space of possible alignments in an attempt to maximize the objective function. Some measures such as sequence similarity or graphlet similarity (Kuchaiev *et al.*, 2010) are defined over pairs of proteins instead of whole alignments. They can be generalized to whole alignments by taking the average similarity score among all pairs of aligned proteins. These measures are called *local measures*, because the contribution of each mapping is independent of the others. On the other hand, global measures aim to evaluate the alignment from a global perspective, and not on a node-to-node basis.

When the objective function is a local measure, the search algorithms used to maximize it usually fall in three categories: greedy best-first algorithms (Neyshabur *et al.*, 2013; Crawford and Milenković, 2014), seed-and-extend algorithms (Kuchaiev *et al.*, 2010; Kuchaiev and Pržulj, 2011; Aladağ and Erten, 2013; Hashemifar and Xu, 2014), and the Hungarian algorithm (Milenković *et al.*, 2010; Crawford and Milenković, 2014). A greedy best-first algorithm starts with an empty alignment and progressively aligns the most similar pair among the available pairs until it is complete. A seed-and-extend algorithm usually starts aligning the most similar node pair (called the seed pair), and then proceeds in a local fashion by aligning the most similar pair among their neighbors. The local nature of seed-and-extend algorithms usually results in large common connected subgraphs. Finally, the Hungarian algorithm can solve the problem optimally as long as the measure optimized is only local. However, optimality in the local objective function does not imply optimality in other measures.

While every method has its own objective function, the alignments are compared according to a set of *target* measures that have been established as the most important (see section 3.1). However, even target measures are of heuristic nature, because except when we align a network with itself, the *correct* mapping is unknown. A good objective function should guide the search algorithm to alignments that score well in all the target measures. Sometimes the objective function can be one of the target measures (Saraph and Milenkovi, 2014). However, since target measures are usually global measures, the search algorithms for local measures described above are in general not applicable.

Among all the methods for aligning PPI networks, currently L-GRAAL (Malod-Dognin and Pržulj, 2015) and HubAlign (Hashemifar and Xu, 2014) seem to yield the best results. All the methods previously mentioned have been shown to be inferior to at least one of them, either directly or indirectly. Both the L-GRAAL and HubAlign objective functions combine a topological measure and sequence similarity. In the case of L-GRAAL, the topological measure is based on graphlets, while its search algorithm uses integer programming. In the case of HubAlign, the topological measure is a local measure that gives higher scores to pairs of topologically important proteins such as those that act as hubs (proteins that have many connections) or bottlenecks. Its search algorithm is a variation of seed-and-extend with multiple seeds.

### 1.3 Our contributions

We present SANA (*Simulated Annealing Network Aligner*), a search algorithm based on Simulated Annealing (Kirkpatrick *et al.*, 1983; Černý, 1985), a metaheuristic local search algorithm with a rich history of successful applications to many optimization problems across a wide variety of domains. It shares many characteristics with MAGNA (Saraph and Milenkovi, 2014), a genetic algorithm. Both can be used to optimize any objective function directly, and both can start with any alignment and then improve it. However, SANA converges to a good solution faster than most existing algorithms. It is conceptually simple to implement, although it has a few user-defined parameters that need to be optimized to get the best results. This can be done either manually by trial-and-error, or automatically (Park and Kim, 1998).

## 2 METHODOLOGY

### 2.1 Main idea

Annealing is a process used in metallurgy to create crystals. A crystal is a highly structured, low-energy state of a material that can only be reached if the material is cooled at a very specific rate. If the substance cools at the "correct" rate, the material can correctly "find" the lowest energy state, which forms the most perfect crystal. It is also similar to a process we're all familiar with: "shaking" a box full of small objects in order to have those objects "settle" in the box. The settled state has lower energy, and is reached by the random motions caused by shaking the box.

Simulated annealing is a metaheuristic algorithm, which means that it is not tailored to any specific optimization problem. It can be applied to any optimization problem as long as the necessary elements are defined: the *solutions*, the *objective function* and the *neighbor relationship*. The analogy to annealing goes as follows: A solution is like an state of the material. In our case, a solution is an alignment. The objective function is analogous to the energy of the material. While in metallurgy the goal is a state with minimum energy, in simulated annealing the best solution minimizes (or maximizes) some arbitrary objective function of our choice. When atoms move due to high temperature, the state of the material changes slightly. In order to simulate this, we need a neighbor relationship that indicates which solutions are close to each other. Then, we can change a solution for a neighbor solution. For instance, we can say two alignments are neighbors if they only vary in one or two mappings of individual pairs of aligned nodes.

If we take a random alignment, called the *initial solution*, we can improve it by looking at its neighbor solutions and choosing the best one. If we repeat this process, we will reach a *local minimum* quickly in which we can no longer improve. However, since the energy landscape is unlikely to be monotonic everywhere, this local minimum is unlikely to be the global minimum. To avoid this pitfall, simulated annealing introduces the ability to allow worse solutions to be selected with some probability, analogous to how high-temperature materials have enough energy to move freely through different states. As the temperature decreases, the ability to escape local minima decreases. If the temperature schedule is chosen correctly, then the solution will tend towards a global minimum.

## 2.2 SANA algorithm

Let $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ be two graphs (networks) with $|V_1| \leq |V_2|$. A pairwise global alignment $a$ (from this point simply alignment) from $G_1$ to $G_2$ is an injective function from $V_1$ to $V_2$. An objective function $f$ is a function from the set of all alignments to the closed range $[0, 1]$.

The basic scheme of SANA is shown in Algorithm 1. The input consists of the two networks $G_1, G_2$, an objective function $f$, a starting alignment (in the absence of one, SANA generates a random alignment), and a maximum execution time $t_{max}$. Since SANA is a generic search algorithm, $f$ can be any objective function. The output is an alignment $a$ that aims to optimize $f$—in our case, we *maximize* various topological and biological similarity measures rather than minimizing energy.

An important element of simulated annealing is the *temperature schedule*, which determines the decline of the probability to accept a worse solution as the algorithm advances. The *temperature* $T(i)$ is a control parameter which depends on the current iteration $i$. It is commonly defined as $T(i) = k \cdot e^{-\lambda \cdot i}$, where $k$ and $\lambda$ are empirically determined constants greater than zero (Kirkpatrick *et al.*, 1983). The temperature is at its highest point at iteration 0 where $T(0) = k$, and approaches 0 asymptotically. The constant $\lambda$ determines how fast the temperature approaches 0.

We define two types of neighbors among alignments: *change* and *swap* neighbors (see Figure 1). Change neighbors differ only in one mapping, which has the same origin in $G_1$ but different destinations in $G_2$. Swap neighbors differ in exactly two mappings, which have the same sources but their images are exchanged. Together, the two types of neighbors allow SANA to explore the solution space completely: through a sequence of neighbors, it is possible to go from any alignment to any other alignment. In SANA when a random neighbor is generated (line 3 of Algorithm 1) the probability of choosing each type of neighbor is proportional to its *branching factor*, i.e., the number of different neighbors of $a$ of that type. This way, all neighbors are equally likely. The idea of using swaps to improve the alignment is not new; it has been used before with other local search algorithms (Chindelevitch *et al.*, 2013; Saraph and Milenkovi, 2014).

In the algorithm, $\Delta E$ denotes the energy increment between the new and the current solution. The probability to accept a worse solution is $P(\Delta E, T(i)) = e^{\Delta E / T(i)}$. This probability decreases when the difference between $f(a)$ and $f(a')$ increases, and it also decreases as the temperature decreases.

As shown in Figure 2, with appropriate values of $k$ and $\lambda$ the temperature schedule allows the algorithm to move freely through the solution space at the beginning, gradually becoming more selective until it stagnates at a local maximum.

## 2.3 Incremental evaluation

A SANA iteration consists of updating the temperature, generating a neighbor alignment, evaluating it, and deciding whether to keep it. It is clear that all the steps other than evaluating the new alignment are constant time operations. Thus, the step that will determine the running time of each iteration is the alignment evaluation. Depending on the objective function, evaluating an alignment may require visiting every node in $G_1$ and $G_2$. However, between two neighbor alignments only one or two mappings are affected. In this section we show that for many typical objective functions, we can efficiently compute the score of the neighbor alignment incrementally from the score of the original alignment.

***Edge Coverage***[1] ***(EC)*** is the fraction of edges in $G_1$ that are aligned to (*i.e.* cover) edges in $G_2$. Coverage represents interactions between

**Algorithm 1** SANA

**input:** $G_1, G_2, f, a_0, t_{max}$
**output:** $a$
1: Let $a = a_0, i = 0$
2: **while** $t_{exec} < t_{max}$ **do**
3: $a' \leftarrow$ random neighbor(a)
4: $\Delta E \leftarrow f(a') - f(a)$
5: **if** $\Delta E \geq 0$ **then** $a \leftarrow a'$
6: **else**
7: $a \leftarrow a'$ with probability $P(\Delta E, T(i))$
8: $i \leftarrow i + 1$
9: **return** $a$

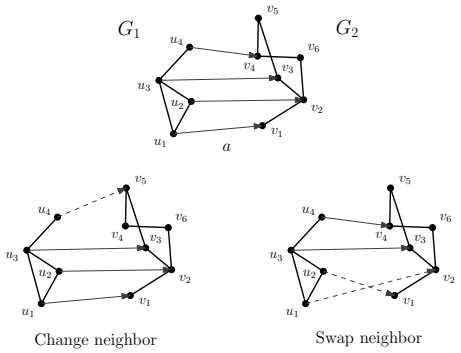


**Fig. 1. Top**: an alignment between $G_1$ on the left and $G_2$ on the right, with the alignment depicted by horizontal arrows. **Bottom left**: a *change* neighbor: the alignment of node $u_4$ has moved from $v_4$ to $v_5$. **Bottom right**: a *swap* neighbor: the alignments $(u_1, v_1)$ and $(u_2, v_2)$ have swapped to be $(u_1, v_2)$ and $(u_2, v_1)$.

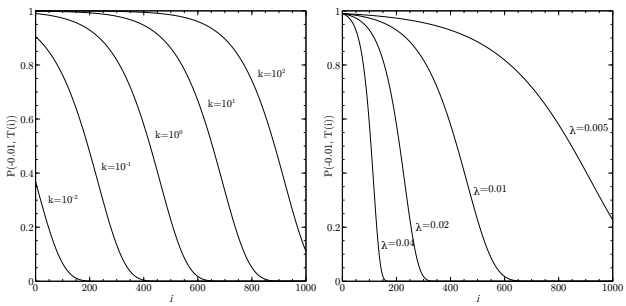


**Fig. 2.** Effect of the $k$ and $\lambda$ constants in the temperature schedule. The figures show the probability to adopt a worse solution with an energy increment $\Delta E = -0.01$ as a function of the iteration $i$. In the left figure, $\lambda$ is fixed to $\lambda = 0.01$. In the right figure $k$ is fixed to $k = 1$.

[1] Often called "Edge Correctness", although coverage is a better term because for most alignments no meaningful definition of "correct" exists. See Appendix.

proteins in one network that are mapped to proteins that also interact in the other, although there is no requirement that the mapped interactions are in any sense biologically equivalent. Let $E_a = \{(u_1, v_1) \in E_1 \mid (a(u_1), a(v_1)) \in E_2\}$ denote the set of edges in $G_1$ that cover edges in $G_2$ in alignment $a$. Then, the edge coverage of $a$ is $EC(a) = |E_a|/|E_1|$.

Evaluating whether an edge is covered requires constant time, so EC can be computed in time $O(|E_1|)$ by checking if each edge in $E_1$ belongs to $E_a$. In order to compute EC incrementally, we need to see how the set of covered edges changes between neighbor solutions. Assume that we know $E_a$ for some alignment $a$ and we need to know $E_{a'}$, where $a'$ is change neighbor of $a$. Let $u_1$ be the node in $G_1$ such that $a(u_1) \neq a'(u_1)$. Note that for any edge not incident to $u_1$, the image of its two endpoints has not changed, and thus whether it is covered by an edge in $G_2$ has not changed either. This means that we only need to check the edges incident to $u_1$. Therefore, the cost of computing $E_{a'}$ incrementally is $O(\text{degree}(u_1))$. The amortized cost of finding $E_{a'}$ for change neighbors is $O(|E_1|/|V_1|)$, the average degree of nodes in $G_1$. The case of swap neighbors is analogous, as we only need to consider the edges incident to the two swapped nodes. For sparse networks such as PPI networks, $|E_1|$ is typically within a small constant factor of $|V_1|$ (see table 3.2) so the practical amortized cost is constant.

***Symmetric Substructure Score (S3)*** (Saraph and Milenkovi, 2014) A drawback of edge coverage is that it does not penalize an alignment for mapping sparse regions of $G_1$ to dense regions in $G_2$. S3 corrects for this by penalizing the mapping of interactions in $G_1$ to non-interactions in $G_2$. Let $\hat{E}_a = \{(u_2, v_2) \in E_2 \mid \exists u_1, v_1 \in V_1 \wedge a(u_1) = u_2 \wedge a(v_1) = v_2\}$ denote the set of edges of the subgraph of $G_2$ induced by the nodes in the alignment $a$. Then, S3 is defined as: $S3(a) = \frac{|E_a|}{|E_1| + |\hat{E}_a| - |E_a|}$.

We have seen how to compute $E_{a'}$ incrementally. The set $\hat{E}_{a'}$ can also be computed similarly. Between change neighbors, if for instance $a(u_1) = u_2$ and $a'(u_1) = v_2$, all the edges incident to $u_2$ will not be in $\hat{E}_{a'}$, and the edges incident to $v_2$ may be in it. Between swap neighbors $a$ and $a''$, we have $\hat{E}_a = \hat{E}_{a''}$. Hence, the cost of computing $\hat{E}_{a'}$ incrementally is $O(\text{degree}(v) + \text{degree}(w))$, and the amortized cost is $O(|E_2|/|V_2|)$.

***Local measures***. Given some notion of similarity between proteins, such as sequence or graphlet similarity, let $s(u, v)$ denote the similarity between two nodes $u \in V_1, v \in V_2$. A local measure $M_s(a) = \sum_{u \in V_1} s(u, a(u))/|V_1|$ is the average similarity of the aligned proteins. We can compute the score difference of a local measure between neighbor alignments by subtracting the similarity of the node pairs that are no longer aligned and adding the similarity of the newly aligned pairs. If we assume that all the required similarities $s(u, v)$ are known before SANA starts, this requires constant time, as neighbor alignments only vary in one or two mappings.

***Weighted Edge Coverage (WEC)*** (Sun *et al.*, 2015) generalizes EC by making the contribution of each covered edge not equal, but a function of the similarity of the aligned endpoints. In order to define the similarity of the nodes, we can use any local measure $s$. Then, for each edge $(u_1, v_1) \in E_1$ covered by the alignment, its contribution is $(s(u_1, a(u_1)) + s(v_1, a(v_1)))/2$. The WEC score is

$$\frac{1}{|E_1|} \sum_{(u_1, v_1) \in E_a} (s(u_1, a(u_1)) + s(v_1, a(v_1)))/2$$

The WEC measure can be optimized incrementally by adapting the method used for EC to account for the weights.

***Combinations*** Any weighted combination of the aforementioned objective functions can also be computed efficiently by computing each one incrementally and then multiplying their scores by the corresponding factor to obtain the final score.

Incremental evaluation makes SANA scalable to enormous PPI networks, as the cost of each step only depends on the average degree of nodes. For objective functions that cannot be evaluated incrementally, SANA may become prohibitively expensive.

# 3 RESULTS

## 3.1 Alignment evaluation

In this section we describe the target measures we use to assess the quality of alignments. These measures are the same used by most previous methods.

For *topological measures*, we use Edge Coverage (EC), Symmetric Substructure Score (S3), Node Correctness (NC) and Largest Common Connected Subgraph (LCCS). The first two are described in Section 2.3.

Assuming there exists a correct alignment between $G_1$ and $G_2$, the *Node Correctness (NC)* is the fraction of correctly aligned nodes. This is the most important measure, but since in general the correct alignment between different species is not known, it can only be used when aligning a network with itself.

The *common subgraph* of an alignment $a$ between $G_1$ and $G_2$ is the subgraph of $G_1$ that remains when considering only edges covered by the alignment: $CS_a = (V_1, E_a)$. A good alignment has a common subgraph with large connected regions. Let $CC_l = (V_l, E_l)$ be the largest connected component of $CS_a$. LCCS measures the size of $CC_l$ as the geometric mean of (i) the fraction of nodes in $CC_l$, $|V_l|/|V_1|$, and (ii) the fraction of edges in $CC_l$, $|E_l|/\min(|E_1|, |\hat{E}_a|)$ (Kuchaiev et al., 2010; more detail in Saraph and Milenković, 2013).

To assess *biological quality*, we measure common Gene Ontology (GO) terms (Ashburner *et al.*, 2000). More precisely, we are interested in the fraction of aligned proteins sharing at least $k$ GO terms, for $k = 1 \ldots 9$. We denote this local measure $\text{GO}_k$. Some GO terms are very common, and thus aligning them does not have much significance. To account for this fact, for each network we removed the most common terms until no more than half of the total GO occurrences remained. Although ostensibly more sophisticated measures using GO terms exist, none has yet become dominant and we believe this simple, easy-to-interpret GO-based measure is the one of the best available at the moment.

## 3.2 Datasets

We attempt a comprehensive analysis by using three different datasets used by several previous studies.

***Noisy yeast*** This dataset consists of six variations of a yeast PPI network (Collins *et al.*, 2007). They all have the same set of nodes and vary only in the number of edges. The first network contains 8,323 interactions, and the other five are formed by progressively increasing the number of edges by 5, 10, 15, 20, and 25% with lower-confidence interactions from the same experiment. We align the first network against itself and each noisy variant. Since the underlying network is the same, we know the true node mapping. This dataset has been used by many previous authors (Saraph and Milenkovi, 2014; Patro and Kingsford, 2012; Sun *et al.*, 2015; Crawford and Milenković, 2014).

***BioGRID dataset*** This dataset includes eight networks from the manually curated BioGRID database (Chatr-aryamontri *et al.*, 2013) (v3.2.101, June 2013). Among the six networks with lowest amount of edges, we align each pair of networks. We also align the two networks with more edges (SC and HS). This dataset was used in Malod-Dognin and Pržulj (2015).

***Yeast and human*** This dataset consists of two older Yeast and Human networks (Collins *et al.*, 2007; Radivojac *et al.*, 2008). This is one of the most often used datasets. It

was used in (Kuchaiev *et al.*, 2010; Neyshabur *et al.*, 2013; Saraph and Milenkovi, 2014; Kuchaiev and Pržulj, 2011).

Note that Yeast and Human have more than one dataset associated with them, for consistent comparison with the results of previous authors.

**Table 1.** List of networks grouped by datasets, with the corresponding identifiers used in the plots.

| Network | Identifier | Proteins | Edges |
|---|---|---|---|
| Yeast | Y0 | 1,004 | 8,323 |
| Yeast (+5% noise) | Y5 | 1,004 | 8,739 |
| Yeast (+10% noise) | Y10 | 1,004 | 9,155 |
| Yeast (+15% noise) | Y15 | 1,004 | 9,571 |
| Yeast (+20% noise) | Y20 | 1,004 | 9,987 |
| Yeast (+25% noise) | Y25 | 1,004 | 10,403 |
| *Rattus norvegicus* | RN | 1,657 | 2,330 |
| *Schizosaccharomyces pombe* | SP | 1,911 | 4,711 |
| *Caenorhabditis elegans* | CE | 3,134 | 5,428 |
| *Mus musculus* | MM | 4,370 | 9,116 |
| *Arabidopsis thaliana* | AT | 5,897 | 13,381 |
| *Drosophila melanogaster* | DM | 7,937 | 34,753 |
| *Saccharomyces cerevisiae* | SC | 5,831 | 77,149 |
| *Homo Sapiens* | HS | 13,276 | 110,528 |
| Yeast | Y | 2,390 | 16,127 |
| Human | H | 9,141 | 41,456 |

### 3.3 Compared methods and parameters

In order to keep our analysis brief and relevant we compare SANA only against the two currently best-performing methods: L-GRAAL (Malod-Dognin and Pržulj, 2015) and HubAlign (Hashemifar and Xu, 2014). Malod-Dognin and Pržulj (2015) performed a thorough analysis demonstrating that for every target measure of interest, either L-GRAAL or HubAlign beat all other methods among GHOST, MI-GRAAL, SPINAL, NETAL, NATALIE, MAGNA, IsoRank, and PISwap (in the BioGRID dataset). Thus, we compare only against HubAlign and L-GRAAL.

We set a maximum execution time of one hour for all methods. In L-GRAAL, we set a limit of 1,000 iterations as in (Malod-Dognin and Pržulj, 2015). We use default values for the remaining parameters of L-GRAAL and the parameters of HubAlign.

For SANA, we need to set the constants $k$ and $\lambda$ of the temperature schedule, the optimal values for which change on a case by case basis depending on variables such as the network sizes or the objective function being evaluated. We use an automatized method to find suitable values automatically in a few minutes. It mimics what humans would do by trial-and-error. As initial temperature $k$, it searches for the lowest temperature such that the behavior at this temperature is still random. For the rate of decay $\lambda$, it searches for the slowest decay such that after $t_{max}$ time the temperature is practically zero. The details of this method are out of scope for this paper.

The objective function of SANA, HubAlign, and L-GRAAL is a combination of sequence similarity and a topological measure. A parameter $\alpha$ controls the weight of sequence similarity as opposed to topological similarity, with $\alpha = 1$ using only sequence, and $\alpha = 0$ using only topology:

$$\text{score}(a) = (1 - \alpha)T(a) + \alpha S(a) \quad (1)$$

where $T(a)$ and $S(a)$ are the topological and sequence similarity of an alignment $a$, respectively. For sequence similarity, we use normalized BLAST bit-scores (Camacho *et al.*, 2009). If $s(u_1, u_2)$ is the BLAST bit-score of proteins $u_1$ and $u_2$,

$$S(a) = \sum_{u_1 \in V_1} \frac{s(u_1, a(u_1))}{\max_{v_1 \in V_1, v_2 \in V_2} s(v_1, v_2)}$$

For the topological measure of SANA, we use one of the following: EC, S3, or the objective function of L-GRAAL, which is WEC with graphlet similarity as local measure (Malod-Dognin and Pržulj, 2015).

Our results suggest that L-GRAAL's objective function provides the best all-round measure of topology for alignment purposes. Moreover, using the same objective function as L-GRAAL illustrates how SANA is the key component in finding the best alignments when combined with any objective function. We denote the resulting method SANA-LG, whereas the variants in which the topological component of the objective function are S3 and EC are denoted SANA-S3 and SANA-EC, respectively.

In the plots, we show the average score of three runs (in Figure 3) or two runs (in the remaining Figures). Moreover, to ensure fairness (specially when aligning a network to itself in the Noisy Yeast dataset), we randomly shuffle the ordering of the nodes and edges in the input network files before each run. Due to this, even deterministic methods such as HubAlign obtain a different score every time.

### 3.4 Topology only comparison

In this section we concentrate on optimizing the topological quality of our alignments without regard to biology (thus $\alpha = 0$). In Figure 3 we see the topological scores obtained with the different methods in the Noisy Yeast dataset, where the true node mapping is known. SANA-LG shows more tolerance to noise, as its scores do not decline as quickly as we add noise. This is important because PPI networks tend to have high levels of noise, as the methods used to discover the interactions have an inherent high error rate.

Note that when $G_1$ and $G_2$ have the same number of nodes, optimizing EC and S3 is equivalent. This is because all the nodes in $G_2$ are part of the alignment, and therefore $\hat{E}_a$, the set of edges in the subgraph of $G_2$ induced by the nodes in the alignment, is simply $E_2$ regardless of the alignment $a$. The only term in the formula of $S3$ that changes as a function of the alignment is $|E_a|$, the number of covered edges, just as in the formula of $EC$. Therefore, for this comparison we omit SANA-S3.

In Figure 4 we see the topological scores of the BioGRID dataset. SANA-LG performs consistently better than L-GRAAL and HubAlign in all measures except in LCCS, and is only beaten in S3 by SANA-S3 (which explicitly optimizes S3) and is only beaten in EC by SANA-EC (which explicitly optimizes EC).

In Figure 5 we compare L-GRAAL's alignment strategy and SANA directly by showing the scores in their common objective function, graphlet-based WEC (with $\alpha = 0$), and see that SANA is capable of optimizing L-GRAAL's objective function better than

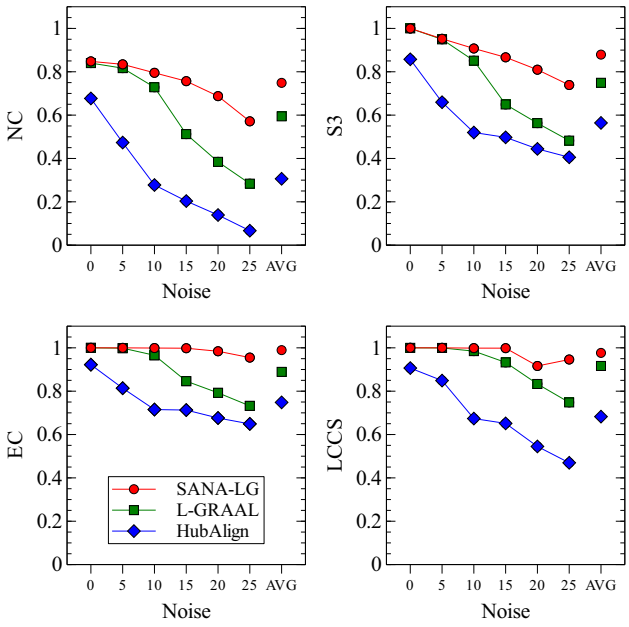


**Fig. 3.** Topological measure scores for the Noisy Yeast dataset when using $\alpha = 0$. $G_1$ is always Y0 (no noise), and the $x$-axis corresponds to the amount of noise in % in $G_2$. For SANA-EC we only show the averages for clarity.

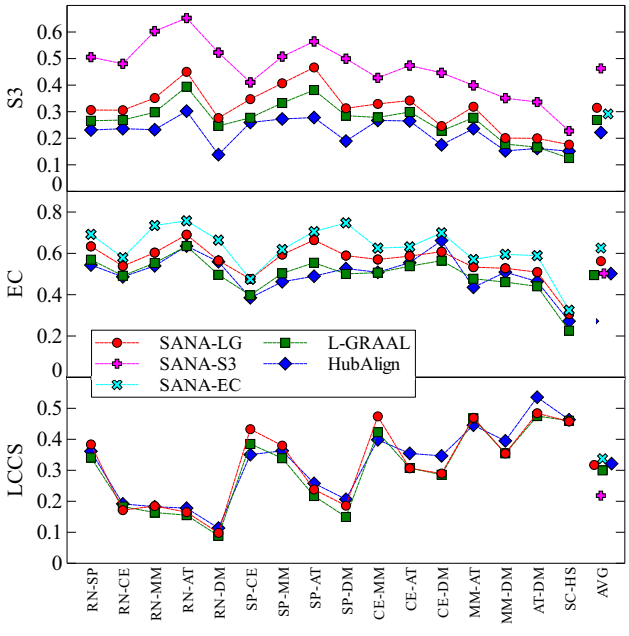


**Fig. 4.** Topological measure scores for the BioGRID dataset when using $\alpha = 0$. Note that the networks are not sorted in any particular order, so the lines are only a visual aid. The top plot shows that SANA-S3 is by far the best method for optimizing S3; similarly, the middle plot demonstrates that SANA-EC is the best at optimizing EC. For SANA-EC and SANA-S3 we only show the averages for clarity, except in their respective objective functions.

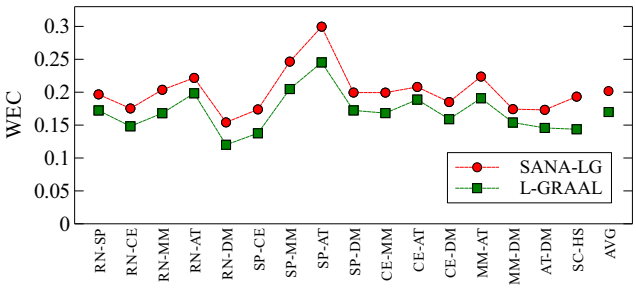


**Fig. 5.** Comparison between L-GRAAL alignment strategy and SANA at optimizing L-GRAAL objective function (graphlet-based WEC) when using $\alpha = 0$. A similar graph appears if we compare SANA-S3 to MAGNA (which optimizes S3).

L-GRAAL itself. We can make similar comparisons with other methods. For instance, in the Yeast and Human dataset, MAGNA, which optimizes S3, obtains a S3 score of $\sim$26% (or $\sim$20% when starting with a random alignment) while SANA-S3 obtains scores of $\sim$40% when starting with a random alignment.

### 3.5 Biological comparison

In this section we evaluate the methods' ability to find alignments that score well in both topology and sequence. It has previously been shown that using sequence information alone to align pairs of nodes in different species produces alignments that have very low topological quality (comparable to random alignments), but that using graphlet-based topological information alone can recover significant biological information (Kuchaiev *et al.*, 2010). Since then it has also been shown that using both types of information can produce alignments that score reasonably well in both the topological and biological sense. However, balancing the two is not as simple as setting $\alpha = 0.5$ in Equation 1 if the topological measure and sequence have different plausible ranges of scores. For instance, when SANA optimizes EC aligning *RNorvegicus* and *SPombe*, it obtains a EC score of 0.69 (Figure 4). However, when SANA optimizes sequence on the same pair of networks, it obtains a Sequence score of 0.057 (Figure 6). This means that using $0.5T(a) + 0.5S(a)$ does not "balance" the two, because sequence scores are under-weighted due to their low value; a relatively small increase in EC would compensate for a large drop in Sequence.

For this reason, we introduce a new balancing parameter called $\beta$ that accounts for differences between expected topological and sequence scores. For instance, $\beta = 0.5$ means that topology and sequence are optimized at equal parts. Given a certain $\beta$, we can compute the $\alpha$ that achieves the corresponding balance as

$$\alpha = \frac{\beta \cdot S_{top}}{(1-\beta) \cdot S_{seq} + \beta \cdot S_{top}}$$

where $S_{top}$ is the score of the method using $\alpha = 0$, and $S_{seq}$ is the score of the method using $\alpha = 1$, i.e. optimizing only sequence[2]. Since $S_{top}$ and $S_{seq}$ vary as a function of the input networks, we compute a different $\alpha$ for each method and for each

[2] HubAlign does not work properly when optimizing Sequence exclusively, so for this method we used $\alpha = 0.9999$ to compute $S_{seq}$.

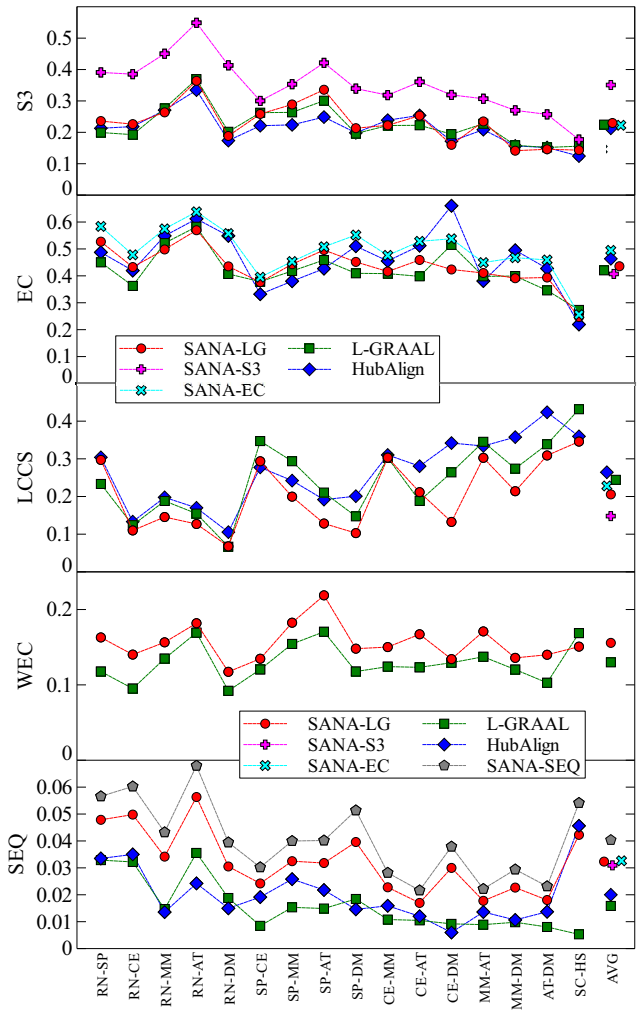


**Fig. 6.** Scores of the three topological target measures, S3, EC, and LCCS, in addition to graphlet-based WEC and Sequence (SEQ), for the BioGRID dataset when setting $\beta = 0.5$. The WEC plot compares L-GRAAL's and SANA-LG's common topological objective function (it is analogous to Figure 5 but with $\beta = 0.5$). For SANA-EC and SANA-S3 we only show the averages for clarity, except in their respective objective functions. In the SEQ plot, we have added SANA-SEQ, which optimizes Sequence alone (i.e., $\alpha = 1$), to give a notion of the upper-bound of Sequence scores for SANA.

network pair. For instance, in the previous example of SANA-EC and *RNorvegicus* and *SPombe*, $\beta = 0.5$ yields $\alpha = 0.92$.

In Figure 6 we see the topological scores of the methods in inter-species alignments when setting $\beta = 0.5$. SANA-LG, HubAlign and L-GRAAL get similar results, except in LCCS, where SANA-LG is generally worse (see discussion below). As expected, SANA-EC and SANA-S3 dominate in their respective objective functions.

Figure 7 shows the biological quality of the alignments obtained by each method. All the variations of SANA outperform the other methods. Thus, SANA is capable of optimizing any topological measure better than the other methods while at the same time obtaining better biological results.

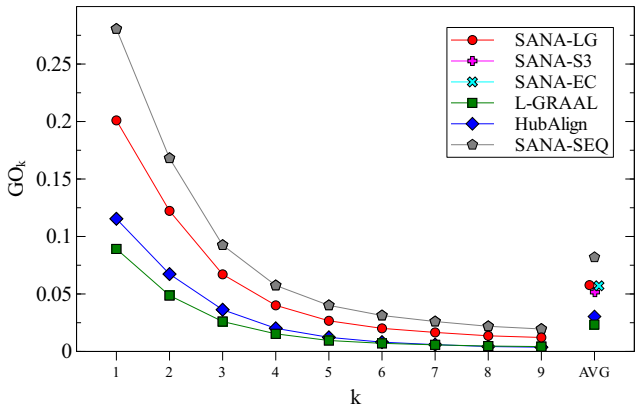


**Fig. 7.** Comparing the biological value of methods: average of each GO$_k$ scores among the 16 network pairs from the BioGRID dataset when using $\beta = 0.5$. For example $k = 1$ measures the fraction of aligned proteins that share 1 GO term; $k = 9$ measures the fraction of aligned proteins that share 9 GO terms. The most egregiously common GO terms are eliminated since they add no useful information (see Section 3.1).

### 3.6 Running time

HubAlign, with a time complexity of $O(n^2 \log n)$ is the fastest method. Its execution time for the largest pair of networks, *SCerevisiae* and *HSapiens*, was 33 minutes. In the case of SANA and L-GRAAL the maximum running time is a parameter specified by the user, which we set to one hour. However, in some instances L-GRAAL finished earlier than that. By design, SANA always uses the full allocated time.

## 4 DISCUSSION

If the past 10 years of network alignment have proven anything, it is that the choice of which objective function(s) is (are) most important is unclear. EC was superseded by S3 because it was noted that EC is too simplistic because it has no incentive to match density fluctuations in the two networks; S3 was designed to explicitly favor alignments that match density fluctuations across the networks. In some sense, EC was important early on because initial EC scores were so low (in the single digits of percentages) that any increase in EC was seen as a good thing. As EC became saturated, it became clear that it was too simplistic and S3 replaced it.

We believe LCCS suffers an analogous shortcoming. When LCCS values were low in the past, pushing them higher was a good thing. But recent results suggest that LCCS values are close to being saturated. Eking out a larger LCCS is not hard today: one could, for example, perform a single "swap" operation (Figure 1) that connects the two largest connected components of the common subgraphs with a single edge, thus creating a larger one, with minimal effect on all other measures. We even considered creating a version of SANA that explicitly optimizes LCCS to prove our point, but unfortunately LCCS cannot easily be incrementally calculated (§2.3), and so our hypothetical SANA-LCCS would take too long to run. At this point, we believe that small differences in LCCS are best sacrificed in favor of better biological significance. This is why we are not concerned that SANA-LG is not best in LCCS; we are happier that all versions

of SANA produce better biological value than any other method while simultaneously producing comparable or better topological results in all measures except LCCS.

It is clear that when SANA explicitly optimizes a particular objective function (EC, S3, or WEC), it does so better than any existing method. What is also interesting is that SANA-LG matches or beats all other methods in these measures while *simultaneously* doing significantly better in biological measures (SEQ and $GO_k$).

While SANA-S3 grossly outperforms all other methods in optimizing S3, its sequence and GO scores are slightly lower than SANA-LG (although SANA-S3 still outperforms all non-SANA methods in sequence and GO measures).

In conclusion, we have introduced a new simulated annealing search algorithm that can optimize any objective function when performing network alignment. We have shown that SANA optimizes any explicit objective function better than any existing method, and produces solutions that score better in virtually every way compared to existing solutions, and does so faster than current methods.

## ACKNOWLEDGMENTS

We are very grateful to Tijana Milenković and Vikram Saraph for directions on how to perform GO similarity analysis. We are also very grateful to Noël Malod-Dognin for assistance with sequence and GO similarity analysis on the BioGRID dataset. We thank Oleksii Kuchaiev for his insights on how to compute a statistically significant measure based on GO terms. We thank Yuki Sawa for help obtaining the sequence data of the Yeast and Human dataset.

*Funding* The research has been supported by the Balsells fellowship.

## APPENDIX: EC IS EDGE COVERAGE, NOT EDGE "CORRECTNESS"

When two pairs of aligned nodes both have an edge between them, we say that the edges are "covered" by each other. Given that there are thousands of edges in both networks, this fact alone, for this particular pair, implies nothing "correct" about this particular alignment. Furthermore, as Figure 3 demonstrates, SAGA-LG is capable of maintaining virtually $100\%$ edge coverage even as the node correctness (and "correctness" here does have meaning) drops significantly.

In other words, there is nothing "correct" about edge correctness. In fact, if one is aligning two cliques, every alignment has $100\%$ EC, even if every node and edge is in the wrong place. The second author (W. Hayes) was in the room when this term was coined for this measure, and at the time it did not occur to anybody that it would be possible that EC could stay so high while the value of the alignment diminished so significantly. It is time to rename this measure to a more meaningful name: edges can be covered independently of their "correctness", and so the word "correct" is misleading. In fact, EC is so incorrect that S3 was invented to replace it. It is time to stop calling it "correct". We are doing a disservice to the long-term community of network alignment scientists by maintaining such a misleading term.

## REFERENCES

Aladağ, A. E. and Erten, C. (2013). Spinal: scalable protein interaction network alignment. *Bioinformatics*, **29**(7), 917–924.

Ashburner, M., Ball, C. A., Blake, J. A., Botstein, D., Butler, H., Cherry, J. M., Davis, A. P., Dolinski, K., Dwight, S. S., Eppig, J. T., Harris, M. A., Hill, D. P., Issel-Tarver, L., Kasarskis, A., Lewis, S., Matese, J. C., Richardson, J. E., Ringwald, M., Rubin, G. M., and Sherlock, G. (2000). Gene Ontology: tool for the unification of biology. *Nature Genetics*, **25**(1), 25–29.

Bayati, M., Gerritsen, M., Gleich, D. F., Saberi, A., and Wang, Y. (2009). Algorithms for large, sparse network alignment problems. In *Proceedings of the 9th IEEE International Conference on Data Mining*, pages 705–710.

Camacho, C., Coulouris, G., Avagyan, V., Ma, N., Papadopoulos, J. S., Bealer, K., and Madden, T. L. (2009). Blast+: architecture and applications. *BMC Bioinformatics*, **10**, 421.

Chatr-aryamontri, A., Breitkreutz, B.-J., Heinicke, S., Boucher, L., Winter, A., Stark, C., Nixon, J., Ramage, L., Kolas, N., ODonnell, L., Reguly, T., Breitkreutz, A., Sellam, A., Chen, D., Chang, C., Rust, J., Livstone, M., Oughtred, R., Dolinski, K., and Tyers, M. (2013). The biogrid interaction database: 2013 update. *Nucleic Acids Research*, **41**(D1), D816–D823.

Chindelevitch, L., Ma, C.-Y., Liao, C.-S., and Berger, B. (2013). Optimizing a global alignment of protein interaction networks. *Bioinformatics*, **29**(21), 2765–2773.

Clark, C. and Kalita, J. (2015). A multiobjective memetic algorithm for ppi network alignment. *Bioinformatics*, **31**(12), 1988–1998.

Collins, S. R., Kemmeren, P., Zhao, X.-C., Greenblatt, J. F., Spencer, F., Holstege, F. C. P., Weissman, J. S., and Krogan, N. J. (2007). Toward a comprehensive atlas of the physical interactome of saccharomyces cerevisiae. *Molecular and Cellular Proteomics*, **6**(3), 439–450.

Cook, S. A. (1971). The complexity of theorem-proving procedures. In *Proceedings of the Third Annual ACM Symposium on Theory of Computing*, STOC '71, pages 151–158, New York, NY, USA. ACM.

Crawford, J. and Milenković, T. (2014). GREAT: GRaphlet Edge-based network AlignmenT. *ArXiv e-prints*.

Davis, D., Yaveroğlu, O. N., Malod-Dognin, N., Stojmirovic, A., and Pržulj, N. (2015). Topology-function conservation in protein-protein interaction networks. *Bioinformatics*.

Elmsallati, A., Clark, C., and Kalita, J. (2015). Global alignment of protein-protein interaction networks: A survey. *Computational Biology and Bioinformatics, IEEE/ACM Transactions on*, **PP**(99), 1–1.

Hashemifar, S. and Xu, J. (2014). HubAlign: an accurate and efficient method for global alignment of proteinprotein interaction networks. *Bioinformatics*, **30**(17), i438–i444.

Hu, J., Kehr, B., and Reinert, K. (2014). NetCoffee: a fast and accurate global alignment approach to identify functionally conserved proteins in multiple networks. *Bioinformatics*, **30**(4), 540–548.

Ito, T., Tashiro, K., Muta, S., Ozawa, R., Chiba, T., Nishizawa, M., Yamamoto, K., Kuhara, S., and Sakaki, Y. (2000). Toward a proteinprotein interaction map of the budding yeast: A comprehensive system to examine two-hybrid interactions in all possible combinations between the yeast proteins. *Proceedings of the National Academy of Sciences*, **97**(3), 1143–1147.

Kelley, B. P., Yuan, B., Lewitter, F., Sharan, R., Stockwell, B. R., and Ideker, T. (2004). Pathblast: a tool for alignment of protein interaction networks. *Nucleic Acids Res*, **32**, 83–88.

Kirkpatrick, S., Gelatt, C. D., and Vecchi, M. P. (1983). Optimization by Simulated Annealing. *Science*, **220**, 671–680.

Klau, G. (2009). A new graph-based method for pairwise global network alignment. *BMC Bioinformatics*, **10**(Suppl 1), S59.

Krogan, N., Cagney, G., Yu, H., Zhong, G., and Guo, X. (2006). Global landscape of protein complexes in the yeast saccharomyces cerevisiae. *Nature*, **440**, 637–643.

Kuchaiev, O. and Pržulj, N. (2011). Integrative network alignment reveals large regions of global network similarity in yeast and human. *BIOINFORMATICS*, **27**, 1390–1396.

Kuchaiev, O., Milenković, T., Memišević, V., Hayes, W., and Pržulj, N. (2010). Topological network alignment uncovers biological function and phylogeny. *Journal of The Royal Society Interface*, **7**(50), 1341–1354.

Li, J., Tang, J., Li, Y., and Luo, Q. (2009). Rimom: A dynamic multistrategy ontology alignment framework. *IEEE Trans. Knowl. Data Eng.*, **21**(8), 1218–1232.

Liao, C.-S., Lu, K., Baym, M., Singh, R., and Berger, B. (2009). Isorankn: spectral methods for global alignment of multiple protein networks. *Bioinformatics*, **25**(12).

Malod-Dognin, N. and Pržulj, N. (2015). L-graal: Lagrangian graphlet-based network aligner. *Bioinformatics*.

Memisevic, V. and Pržulj, N. (2012). C-graal: Common-neighbors-based global graph alignment of biological networks. *Integr. Biol.*, **4**, 734–743.

Milenković, T., Ng, W. L., Hayes, W., and Pržulj, N. (2010). Optimal network alignment with graphlet degree vectors. *Cancer Informatics*, **9**, 121–137.

Milenković, T., Zhao, H., and Faisal, F. E. (2013). Global network alignment in the context of aging. In *Proceedings of the International Conference on Bioinformatics, Computational Biology and Biomedical Informatics*, BCB'13, pages 23:23–23:32, New York, NY, USA. ACM.

Neyshabur, B., Khadem, A., Hashemifar, S., and Arab, S. S. (2013). Netal: a new graph-based method for global alignment of proteinprotein interaction networks. *Bioinformatics*, **29**(13), 1654–1662.

Park, M.-W. and Kim, Y.-D. (1998). A systematic procedure for setting parameters in simulated annealing algorithms. *Computers and Operations Research*, **25**(3), 207 – 217.

Patro, R. and Kingsford, C. (2012). Global network alignment using multiscale spectral signatures. *Bioinformatics*, **28**(23), 3105–3114.

Pržulj, N., Corneil, D. G., and Jurisica, I. (2004). Modeling interactome: scale-free or geometric? *Bioinformatics*, **20**(18), 3508–3515.

Radivojac, P., Peng, K., Clark, W. T., Peters, B. J., Mohan, A., Boyle, S. M., and Mooney, S. D. (2008). An integrated approach to inferring genedisease associations in humans. *Proteins: Structure, Function, and Bioinformatics*, **72**(3), 1030–1037.

Saraph, V. and Milenkovi, T. (2014). Magna: Maximizing accuracy in global network alignment. *Bioinformatics*.

Singh, R., Xu, J., and Berger, B. (2008). Global alignment of multiple protein interaction networks with application to functional orthology detection. *Proceedings of the National Academy of Sciences*, **105**(35), 12763–12768.

Sun, Y., Crawford, J., Tang, J., and Milenkovi, T. (2015). Simultaneous optimization of both node and edge conservation in network alignment via wave. In M. Pop and H. Touzet, editors, *Algorithms in Bioinformatics*, volume 9289 of *Lecture Notes in Computer Science*, pages 16–39. Springer Berlin Heidelberg.

Uetz, P., Dong, Y.-A., Zeretzke, C., Atzler, C., Baiker, A., Berger, B., Rajagopala, S. V., Roupelieva, M., Rose, D., Fossum, E., and Haas, J. (2006). Herpesviral protein networks and their interaction with the human proteome. *Science*, **311**(5758), 239–42.

Černý, V. (1985). Thermodynamical approach to the traveling salesman problem: An efficient simulation algorithm. *Journal of Optimization Theory and Applications*, **45**(1), 41–51.

Vijayan, V., Saraph, V., and Milenković, T. (2015). Magna++: Maximizing accuracy in global network alignment via both node and edge conservation. *Bioinformatics*.

Zaslavskiy, M., Bach, F., and Vert, J.-P. (2009). A path following algorithm for the graph matching problem. *IEEE Trans. Pattern Anal. Mach. Intell.*, **31**(12), 2227–2242.

Zhang, Y. and Tang, J. (2013). Social network integration: Towards constructing the social graph. *CoRR*, **abs/1311.2670**.